

## ARCHITECTURE FOR VIRTUAL PRIVATE NETWORKS

### Inventors:

Henk J. Bots  
William E. Hunt  
Derek Palma  
John Lawler

This is a continuation of Application No. 08/874,090,  
filed June 12, 1997.

## BACKGROUND OF THE INVENTION

### 1. Related Information

The present invention is related to the one described  
in copending U.S. Patent Application entitled "An Apparatus for  
Implementing Virtual Private Networks," Serial No. 08/874,091  
[Attorney Docket No. 20155-702], assigned to the assignee of the  
present application and filed concurrently herewith.

### 2. Field of the Invention

The present invention relates to the field of data  
communications. More particularly, the present invention relates  
to techniques for implementing secure virtual private networks

over public or otherwise insecure data communications infrastructures.

### 3. Background

In recent years organizations have come to rely heavily on the ability to transmit electronic data between members of the organization. Such data typically includes electronic mail and file sharing or file transfer. In a centralized, single site organization, these transfers of electronic data are most commonly facilitated by a local area network (LAN) installed and operated by the particular enterprise.

Preventing unauthorized access to data traversing an enterprise's LAN is relatively straightforward. This applies to both unauthorized accesses by members of the enterprise and, more importantly, to third parties on the outside. As long as intelligent network management is maintained, unauthorized accesses to data traversing an enterprise's internal LAN are relatively easily avoided. It is when the enterprise spans multiple sites that security threats from the outside become a major concern.

For distributed enterprises that desire the conveniences of the above-described electronic data transfers, there are several options that exist today, but each with associated disadvantages. The first option is to interconnect the offices or various sites with dedicated, or private communications connections often referred to as leased lines. This is the traditional method

organizations use to implement a wide area network (WAN). The disadvantages of implementing an enterprise owned and controlled WAN are obvious: they are expensive, cumbersome and frequently underutilized if they are established to handle the peak capacity requirements of the enterprise. The obvious advantage to this approach is that the lines are dedicated for use by the enterprise and are therefore secure, or reasonably secure, from eavesdropping or tampering by intermediate third parties.

An alternative to the use of dedicated communications lines in a wide area network is for an enterprise to handle intersite data distributions over the emerging public network space. Over recent years, the Internet has transitioned from being primarily a tool for scientists and academics to a mechanism for global communications with broad ranging business implications. The Internet provides electronic communications paths between millions of computers by interconnecting the various networks upon which those computers reside. It has become commonplace, even routine, for enterprises, even those in nontechnical fields, to provide Internet access to at least some portion of the computers within the enterprise. For many businesses this facilitates communications with customers, potential business partners as well as the distributed members of the organization.

Distributed enterprises have found that the Internet is a convenient tool to provide electronic communications between members of the enterprise. For example, two remote sites within

the enterprise may each connect to the Internet through a local Internet Service Provider (ISP). This enables the various members of the enterprise to communicate with other sites on the Internet including those within their own organization. The

5 limiting disadvantage of using the Internet for intra-enterprise communications is that the Internet is a public network space. The route by which data communication travel from point to point can vary on a per packet basis, and is essentially indeterminate. Further, the data protocols for transmitting information over the

10 various networks of the Internet are widely known, and leave electronic communications susceptible to interception and eavesdropping with packets being replicated at most intermediate hops. An even greater concern arises when it is realized that communications can be modified in transit or even initiated by

15 impostors. With these disconcerting risks, most enterprises are unwilling to subject their proprietary and confidential internal communications to the exposure of the public network space. For many organizations it is common today to not only have Internet access provided at each site, but also to maintain the existing

20 dedicated communications paths for internal enterprise communications, with all of the attendant disadvantages described above.

While various encryption and other protection mechanisms have been developed for data communications, none completely and

25 adequately addresses the concerns raised for allowing an

enterprise to truly rely on the public network space for secure intra-enterprise data communications. It would be desirable, and is therefore an object of the present invention to provide such mechanisms which would allow the distributed enterprise to rely

5 solely on the public network space for intra-enterprise communications without concern for security risks that presently exist.

## SUMMARY OF THE INVENTION

From the foregoing it can be seen that it would be desirable and advantageous to develop protocols and architecture to allow a single organization or enterprise to rely on the public network space for secure intraorganizational electronic data communications. The present invention is thus directed toward the protocols and architecture for implementing secure virtual private networks over the Internet or other public network systems. The architecture of the present invention introduces a site protector or virtual private network (VPN) unit which moderates data communications between members of a defined VPN group. In accordance with one embodiment of the present invention, the site protector resides on the WAN side of the site's router or routing apparatus which is used to connect the enterprise site to the Internet. In alternative embodiments, the site protector will reside on the LAN side of the router. The essential point for all embodiments is that the site protector be in the path of all relevant data traffic.

To ensure secure data communications between members of the same VPN group, the site protector or VPN unit implements a combination of techniques for data packet handling when packets are to be sent between members of the group. The packet handling processes include various combinations of compression, encryption and authentication, the rules for each of which may vary for members of different groups. For each group defined as a virtual

private network, the various parameters defining the compression, encryption and authentication are maintained in lookup tables in the associated VPN units. The lookup tables maintain information not only for fixed address members of the group but support is

5 also provided for remote clients. This ability allows remote users to dial into a local Internet Service Provider and still maintain membership in a virtual private network group for secure communications over the Internet with other members of the group. In the case of a remote client, the site protector may, in one

10 embodiment, be simulated by software running on the remote client.

In other aspects of the present invention, the VPN units or site protectors may be dynamically configured to add or subtract members from the virtual private network group or

15 recognize their movement, or change other parameters affecting the group. Various other packet handling aspects of the invention include addressing the problem of some data packets growing too large by the inclusion of encryption and authentication information. Another packet handling aspect

20 provides a mechanism for Internet communications which hides information identifying the source and destination of the data packet. In this aspect of the present invention, the VPN units are treated as the source and destination for the Internet communication data packets with the VPN units encapsulating the

25 source and destination addresses of the endstations.

## BRIEF DESCRIPTION OF THE DRAWINGS

The objects, features and advantages of the present invention will be apparent from the following detailed description, in which:

5           Figure 1 illustrates a prior art configuration for an exemplary enterprise's intraenterprise communication architecture.

10           Figure 2 illustrates an enterprise communication scenario in accordance with the present invention utilizing the Internet or other public network space as the vehicle for conveying messages between members of a virtual private network.

15           Figure 3 illustrates a flow diagram for the handling of a packet being transmitted from one member of a virtual private network group to another member over the Internet.

20           Figure 4 illustrates the handling of a data packet received over the Internet by one member of a virtual private network group from another member.

          Figure 5 illustrates graphically the life cycle of a data packet being sent from one member of a virtual private network group to another over the Internet.

          Figure 6 illustrates an alternate life cycle of a data packet being sent from one member of a virtual private network



group to another over the Internet where the source and destination addresses of the group members are also concealed.

005077 7509760

## DETAILED DESCRIPTION OF THE INVENTION

Protocols and an architecture are disclosed for implementing secure virtual private networks for enterprise communications over the Internet or other public network space.

5 Although the present invention is described predominantly in terms of utilizing the Internet as a communications medium, the concepts and methods are broad enough to accomplish the implementation of secure virtual private networks over other public or insecure communications media. Throughout this

10 detailed description, numerous specific details are set forth such as particular encryption or key management protocols, in order to provide a thorough understanding of the present invention. To one skilled in the art, however, it will be understood that the present invention may be practiced without

15 such specific details. In other instances, well-known control structures and system components have not been shown in detail in order not to obscure the present invention.

In many instances, components implemented by the present invention are described at an architectural, functional

20 level. Many of the elements may be configured using well-known structures, particularly those designated as relating to various compression or encryption techniques. Additionally, for logic to be included within the system of the present invention, functionality and flow diagrams are described in such a manner

that those of ordinary skill in the art will be able to implement the particular methods without undue experimentation. It should also be understood that the techniques of the present invention may be implemented using a variety of technologies. For example,

5 the virtual private network unit or site protector to be described further herein may be implemented in software running on a computer system, or implemented in hardware utilizing either a combination of microprocessors or other specially designed application specific integrated circuits, programmable logic

10 devices, or various combinations thereof. It will be understood by those skilled in the art that the present invention is not limited to any one particular implementation technique and those of ordinary skill in the art, once the functionality to be carried out by such components is described, will be able to

15 implement the invention with various technologies without undue experimentation.

Referring now to Figure 1 there is shown a traditional scenario for intra-enterprise data communications for a distributed organization. In this illustration of an exemplary

20 organization configuration, the enterprise consists of a headquarters location 105 with additional sites or branches 110 and 112, respectively. In modern organizations, such as the exemplary one of Figure 1, the headquarters' site 105 as well as the branch sites 110 and 112 may each comprise numerous

25 personnel, many of whom are provided with computers or work

stations with network access. The internal network configurations at the headquarters for branches may take many forms including one or several local area networks (LANs). For intersite communications between headquarters and the branches, 5 dedicated or leased communications lines 115 and 120 may be provided. In addition, an optional dedicated communications path 125 may be provided between the branches 110 and 112. As an alternative to the optional dedicated communications line 125 between the branches, data packets between branch 110 and branch 10 112 may be routed through the headquarters' network equipment.

09740694 "40690  
In addition to the dedicated communications lines between the headquarters and the various branches, it is common today to provide computer users within an organization access to the Internet for electronic mail to external parties as well as 15 for doing various types of research over the Internet using such tools as the World Wide Web, etc. As shown in Figure 1, the usual scenario where the headquarters' site 105 and the branches 110 and 112 are each separately provided with direct access to Internet Service Providers 130, 133 and 136, respectively. This 20 facilities the users at the various sites with their access to the Internet for the above purposes. In an alternate configuration, it may be that only the headquarters site 105 is provided with access to an Internet service provider 130 and that users of the computers of the branch sites 110 and 112 will 25 connect to the Internet through headquarters via their dedicated

communications paths 115 and 120. The downside to this alternate configuration is that it greatly increases the bandwidth utilization on the dedicated lines, perhaps to the point of saturation. An advantage is that only one gateway to the Internet need be provided for the organization which simplifies enforcing security constraints on connections to the outside world.

In the exemplary organization 100, it is also shown that in some circumstances it may be desirable to allow customers or other business partners to dial in directly to the computer network of the organization. In Figure 1 it is illustrated that the customer 140 may in fact carry out such communications over a communications path 145 which may be a dedicated line provided between the customer and the organization for the customer's convenience. The path 145 may also be a dial-up line which the customer might use only sporadically. Consistent with the emerging use of the Internet and its popularity, the customer 140 is shown having its own Internet connection through ISP 148.

Finally, there is shown in Figure 1 that it is frequently desirable for other members of the enterprise who may be on the road or working from home or other remote locations to exchange data with other members of the enterprise. There is thus shown remote clients 150 and 155 communicating with the headquarters over long distance telephone lines 157 and 158.

This example assumes that the remote clients are in a truly remote location from the headquarters. The remote clients 150 and 155 are also respectively shown having local access to the Internet through local ISPs 160 and 165.

5           The above description of an enterprises data communications configuration according to Figure 1 illustrates the disadvantages described in the previous section. These disadvantages are eliminated by implementation of the present invention as illustrated generally with reference to Figure 2.

10 In the enterprise network communication configuration 200 illustrated in Figure 2, the headquarters 105, first branch 110 and second branch 112 of the organization are illustrated in a more detailed logical way then presented in Figure 1. Thus, the headquarters 105 is illustrated with three endstations 201, 202 and 203, respectively coupled to communicate data packets over

15 local area network (LAN) 205. Likewise, the branch site 110 is shown having a plurality of endstations 211, 212 and 213. respectively coupled to communicate data locally over LAN 215. Finally, the second branch site 112 is shown with an illustrative

20 set of computer stations 221, 222 and 223 connected to communicate over LAN 225. The customer site 140 is also illustrated in Figure 2 as comprising of plurality of computers illustrated by 331 and 332 coupled to communicate over the customer's LAN 235. The local area networks utilized for data

25 communications within the headquarters, customer and branch sites

may adhere to a wide variety of network protocols, the most common of which are Ethernet and Token Ring.

As can be seen in Figure 2, the dedicated communications lines between the headquarters site 105 and the branch sites 110 and 112 as well as between the headquarters site 105 and the customers site 140 have been eliminated. Instead, in accordance with the present invention data communications between members of the organization are intended to be carried out over the Internet or other public network space. For purposes of the present invention, it will be assumed that it is the widely emerging Internet that will be the medium for data packet transfers between members of the organization.

Each of the LANs for the particular sites illustrated in Figure 2 ultimately interconnect to the Internet 250 through an associated routing or gateway device which are identified as routers 240, 242, 244 and 246, respectively. It is to be understood that data packets conveyed between a various sites illustrated in 200 would traverse, in many cases, a plurality of additional routing devices on their way between the source and destination sites for the packets. The mechanisms for data packet transfers over the Internet are well known and are not described in great detail herein. It is understood that data packets are assembled in accordance with the Internet Protocol (IP) and are referred to herein as IP packets regardless of the

version of the Internet protocol presently in effect. In the case of the remote clients 150 and 155 illustrated in Figure 2 it is understood that they utilize communication software to dial up a local Internet service provider which itself provides the gateways necessary for communications over the Internet 250.

As has been described above, prior efforts to utilize the Internet for secure data communications have required an awareness or implementation of security considerations at the endstations. This is disadvantageous when transparency to an end user is desirable. The present invention, on the other hand is transparent to end users with data communications over the Internet occurring exactly as they appear to have before. However, for users identified as members of the same virtual private network, data communications are handled in a manner that assures the security and integrity of the data packets. Illustrated in Figure 2, between the Internet 250 and each of the respective routers 240, 242, 244 and 246, are Virtual Private Network Units (VPNUs) 250, 252, 254 and 256. In accordance with the particular illustrated embodiment of the present invention, the VPNUs reside between a site's router and the path to the Internet. It should be understood that this placement of VPN units in the overall system architecture represents only one placement choice. It will be clear from the materials that follow that the key point with respect to VPNU placement is that they reside in the path of data traffic. In many embodiments, it



may in fact prove desirable to situate the VPNU on the LAN side of a site's router. As will be described in more detail below, the VPN units maintain lookup tables for identifying members of specific virtual private network groups.

5           When a data packet is sent between source and destination addresses that are both members of the same VPN group, the VPNU will process the data packet from the sending side in such a way as to ensure that it encrypted, authenticated and optionally compressed. Likewise, the VPNU servicing the site  
10 where the destination address is located will detect that a packet is being propagated between members of the same VPN group. The receiving VPNU will handle the process of decrypting and authenticating the packet before forwarding it toward the destination endstation. In this way, secure data communications  
15 between end users is effected in a manner that is transparent to the end users. In the case of remote clients 150 and 155, the VPNU may be simulated in software which operates in conjunction with the communication software for connecting the remote client to the associated local Internet service provider.

20           The functionality of the VPN units will be described with reference to the following figures beginning with the flowchart of Figure 3. When a data packet originates from an endstation, such as endstation 202 of LAN 205 at site 105, and its destination is to a remote site, other than the headquarters

site 105, it will initially be treated as an ordinary Internet data packet transfer. The packet will proceed from the endstation 202 over the LAN 205 to the routing device 240 which will encapsulate the data packet in accordance with the Internet Protocol, forming an outbound IP packet. On its way out of the site, the IP packet will pass through the associated VPN unit for the site. The flowchart illustrated at Figure 3 shows the functional operation of a VPN unit for an outbound packet that is received thereby. The Transmit Packet procedure 300 begins when the outbound data packet is received at the VPNU at step 310. At decision box 320, it is determined whether or not the source and destination addresses for the data packet are both members of the same VPN group. This determination may be made with reference to lookup tables that are maintained by the VPN units or reference to other memory mechanisms. This step may be thought of as member filtering for data packets being transmitted between the particular site and the VPN unit which services it. If the source and destination address for the data packet are not both members of the same VPN group, then at step 330 the packet is forwarded to the Internet as ordinary Internet traffic from the site as though the VPNU were not involved. In which case, the procedure ends at step 335. In one alternative embodiment, it may be desirable to discard data traffic that is not destined between members of a VPN group rather than forwarding it as unsecure traffic. In another alternative embodiment, it may be

desirable to provide the option to either pass or discard non-VPN-group data traffic.

If, at decision box 320, the member filter, it is determined that both the source and destination addresses for the data packet are members of the same VPN group, then the data packet is processed at step 340 undergoing various combinations of compression, encryption and authentication. The lookup tables maintained by the VPN unit 250 and all of the VPN units, in addition to identifying members of particular VPN groups, also identify whether or not data packets transferred between members of the particular VPN group are to be compressed and if so, what algorithm is to be used for compression. Many possible compression algorithms are well-known, but in one embodiment of the invention, LZW compression is implemented. The lookup table for the VPN group of which the source and destination addresses are members also identifies the particular encryption algorithm to be used for data packets traversing the Internet for that VPN group as well as the authentication and key management protocol information to be used thereby. As an alternative to lookup tables, the VPNU may be programmed to always use the same algorithms for all VPN groups.

The particular packet processing algorithms to be used for VPN traffic may vary, so long as the lookup tables in both the sending and receiving VPN units identify the same

compression, encryption and authentication rules and are capable of implementing and deimplementing them for members of the same group. It is to be understood that a single VPNU may serve multiple VPN groups and that particular addresses may be members of multiple groups. Thus, at step 340, when a packet is destined from one member of the VPN group to another, the packet is processed according to the compression, encryption and authentication rules identified in the VPNU tables for that particular VPN group. Then, at step 350, the processed packet is forwarded toward the destination address over the Internet. The procedure of the sending VPN unit then ends at step 355.

The receiving VPNU reverses the above processes for VPN traffic as illustrated by the flowchart of Figure 4. The Receive Packet procedure 400 begins at step 410 when an inbound data packet is received from the Internet at the receiving VPN unit. At decision box 420, the inbound data packet is examined to determine if the source and destination addresses of the data packet are both members of the same VPN group. It is assumed that the lookup tables maintained by all of the VPN units are both consistent and coherent. If the inbound data packet is determined not to be VPN traffic, then the packet is passed through and forwarded to the receiving site as though it were normal Internet data traffic at step 430. In which case the process ends at step 435. In one alternative embodiment, it may

be desirable to discard incoming data traffic that is not from an identified member of a VPN group supported by the VPNU.

For data packets that are determined to be VPN traffic at decision box 420, the VPN unit will process the inbound packet to recover the original data packet as it was provided from the source endstation. The lookup table maintained by the receiving VPN unit will identify the compression, encryption and authentication rules used for the VPN group and reconstruct the original IP packet in accordance with those rules at step 440.

Then, the reconstructed packet will be delivered to the site of the destination address at 450 with the procedure ending at step 455.

Figure 5 illustrates graphically the life cycle of the data packet sent between two members of the same VPN group. The data packet originates from a source 500 and propagates from the sources site through its associated router to generate IP data packet 510. The data packet 510 is not intended to illustrate all the fields associated with a complete IP data packet, but shows the relevant portions for this discussion which include the destination address, source address and the payload information of the packet. The data packet 510 is then examined by the VPN unit which determines whether the data packet is traffic between members of an identified VPN group. The VPN unit 520 processes the packet in accordance with the packet processing procedures

described above with respect to Figure 3 with the resulting packet being illustrated as packet 530. Packet 530 still identifies the destination and source addresses of the data packet, but the remainder of the packet is encrypted, and  
5 optionally compressed.

Following processing by the outbound VPNU, the data packet is propagated through the Internet to 550 with the destination and source information identifying to the associated routers of the Internet the path by which the packet should  
10 ultimately take to reach its destination. The packet emerges from the Internet at the edge of the destination site as data packet 540 which is essentially identical to the data packet 530. The packet is "deprocessed" by the receiving VPN unit 550 which restores the original packet into its form 560 for delivery to  
15 the ultimate destination through the receiving site's associated router at destination 570.

As was described above, the present invention approach to virtual private networks supports not only optional compression of data packets, but encryption and authentication  
20 techniques as well. One emerging standard for key management in connection with Internet Protocol data transfers with authentication is referred to as simple key management for Internet Protocol (SKIP) which is described by US Patent 5,588,060 assigned to Sun Microsystems, Inc. of Mountain View,

CA. Authenticated data transfers using SKIP support a mode of data transfer referred to as tunnel mode. The above described data transfer with respect to Figure 5 illustrates a transport mode of operation in which the data and source addresses are exposed as the data packet traverses the Internet. In tunnel mode, an added measure of security may be provided by encapsulating the entire data packet in another packet which identifies the source and destination addresses only for the VPN units. This conceals the ultimate source and destination addresses in transit.

Figure 6 illustrates the life cycle of a data packet being propagated from a source 600 to a destination 670 utilizing tunnel mode. In this mode of operation, the data packet 610 is processed by outbound VPNU 620 which generates a resulting packet 630. The resulting packet 630 encrypts and compresses (optionally) not only the data payload of the packet, but the destination and source addresses of the endstations as well. The encapsulated packet is then provided with an additional header that identifies that the source of the packet is the outbound VPNU 620 and that the destination is the inbound VPNU 650. Thus, the packet 640 which emerges from the Internet is identical to the packet 630 with respect to its source and address information and encapsulated payload. The packet is decomposed by the inbound VPNU 650 to reconstruct the original data packet at 660 for delivery to the destination 670.

005097-140900  
The overall architecture of the present invention is robust. It allows end users the convenience of proprietary data communications to take place over a public network space such as the Internet. The architecture of the present invention also  
5 allows a wide variety of compression, encryption and authentication technologies to be implemented, so long as the VPN units at each end of the transaction support the associated protocols. The present invention is also capable of working in concert with traditional Internet security mechanisms such as  
10 corporate firewalls. A firewall might operate in series with the VPN unit at a given site, or, intelligently be configured in a single box with the VPN unit to provide parallel firewall and VPN unit security functions.

There has thus been described a protocol and  
15 architecture for implementing virtual private networks for using a public network space for secure private network data communications. Although the present invention has been described with respect to certain exemplary and implemented embodiments, it should be understood that those of ordinary skill  
20 in the art will readily appreciate various alternatives to the present invention. Accordingly, the spirit and scope of the present invention should be measured by the terms of the claims which follow.